LA-UR-98-2510

Reconfigurable Computer Array: The Bridge between High Speed Sensors and Low Speed Computing

Scott H. Robinson[1], Michael P. Caffrey[2], and Mark E. Dunham[1]

[1]Los Alamos National Laboratory, Space Engineering, NIS-4 MS-D448,
Los Alamos, New Mexico USA 87545
{shr, mdunham}@lanl.gov
[2]Los Alamos National Laboratory, Space Data Systems, NIS-3 MS-D440,
Los Alamos, New Mexico USA 87545
mpc@lanl.gov

Abstract. A universal limitation of RF and imaging front-end sensors is that they easily produce data at a higher rate than any general-purpose computer can continuously handle. Therefore, Los Alamos National Laboratory has developed a custom Reconfigurable Computing Array board to support a large variety of processing applications including wideband RF signals, LIDAR and multi-dimensional imaging. The boards design exploits three key features to achieve its performance. First, there are large banks of fast memory dedicated to each reconfigurable processor and also shared between pairs of processors. Second, there are dedicated data paths between processors, and from a processor to flexible I/O interfaces. Third, the design provides the ability to link multiple boards into a serial and/or parallel structure.

1 Computational Challenges

In modern digital remote sensing applications the need to continuously process enormous amounts of data is always present. Front-end sensors can easily generate a 100 megabytes per second or more of raw data. The ultimate goal is to reduce massive amounts of raw data down to answers that are useful and can be reasonably understood, stored or transmitted. Our challenge is to develop systems that continuously process data efficiently, are flexible to changing requirements and do it all at a reasonable system cost.

Los Alamos National Laboratory1 (LANL) has a number of these types of computational challenges. As part of the Laboratory's research into lightning, RF propagation and astrophysics, there is a strong requirement for processing wideband RF data. Analog-to-digital converters running at rates of 100M samples per second and more support this work. Currently, snapshots of RF data are recorded, relayed back to the

[1] The University of California operates Los Alamos National Laboratory for the United States Department of Energy under contract W-7405-ENG-36. Publication LA-UR-98-2510.

laboratory and analyzed. Before system performance can improve significantly, the analysis of this data must move closer to the sensor and occur continuously.

Another effort with a great need for real-time signal processing at the Laboratory is a large and ongoing LIght Detection And Ranging (LIDAR) research program. The needs include: determining the range to a target when using a laser and detector operating from a moving platform; data filtering and signal averaging to increase the signal-to-noise ratio; return pulse integration and pattern recognition. Current upgrades will increase the laser-firing rate, add multiple laser frequencies, incorporate pointing information and expand the catalog of recognized return signals. In order for this work to continue, LANL must improve the signal processing hardware of the system beyond what is currently available from general purpose computers and DSP's.

Multi-dimensional imaging is another area of intense research at the Laboratory and has its own computational challenges. Multi-dimensional image data are traditionally collected in the form of datacubes with spatial, spectral and/or time information on each axis. Analysis of these datacubes requires the application of complex image processing algorithms. Again, generally available hardware is incapable of providing the data throughput required by these systems.

All of these efforts put severe demands on the signal processing hardware and data throughput requirements. For proper system operation it is critical that signal processing takes place continuously at the full rate of the incoming data. However, these problems share common characteristics that can be mutually addressed. A key to processing is the fact that the data has dependencies that are block oriented. Although these blocks can be large, they are relatively independent of each other. This allows pipelined operations to be performed on blocks of data either in serial or parallel pipe configurations. Related to this is the system characteristic that data latency is not critical. Therefore, additional pipeline stages can be added to allow throughput to increase or to provide other system functionality without affecting overall computational performance.

2 LANL Developed RCA Board

To meet the signal processing challenge of the systems described, the Space Engineering group and several Space Science groups at the Laboratory have joined together to develop a Reconfigurable Computer Array (RCA) board. The RCA-2 board bridges the gap between high speed front-end ASIC's and low speed back-end general purpose computers. The board was specifically designed to process large amounts of data with block oriented dependencies at high throughput rates. Additionally, the RCA-2 board can be reconfigured to meet the needs of individual systems and can deliver the required computational horsepower at a reasonable cost.

The goals of the project were to design a board that: 1) could support continuous data throughput of at least 100 Mbytes per second; 2) could provide the ability to link multiple boards together into a serial and/or parallel processing structure; and 3) could be configured to a wide range of signal and image processing requirements.

These goals translated into three key design features. First, around each reconfigurable processor the design provides a large amount of fast memory to support block oriented processing without resource contentions. Second, the design provides dedicated data paths between processors, and from a processor to a flexible I/O interface. These wide data communication channels keep blocks of data moving continuously through the board. Finally, the design provides three flexible, high speed input/output ports to allow data to move between boards as required by the processing algorithms. These I/O ports are accessible at the front panel. The RCA-2 board uses small daughter cards to adapt to different input or output data formats. With three front panel ports, two separate input data streams can be processed together to produce one output data stream. Alternatively, one stream of input data can be processed to create two streams of output data.

3 Function Description

Figure #1 shows the RCA-2 board and labels the major components. The design is implemented on a 'C' size VXI card that measures 233mm in width by 340mm in length. Attached to the board in daughter card location #1 is an input card. An output card is attached at daughter card location #3. Daughter cards measure 78mm in width by 142mm in length.

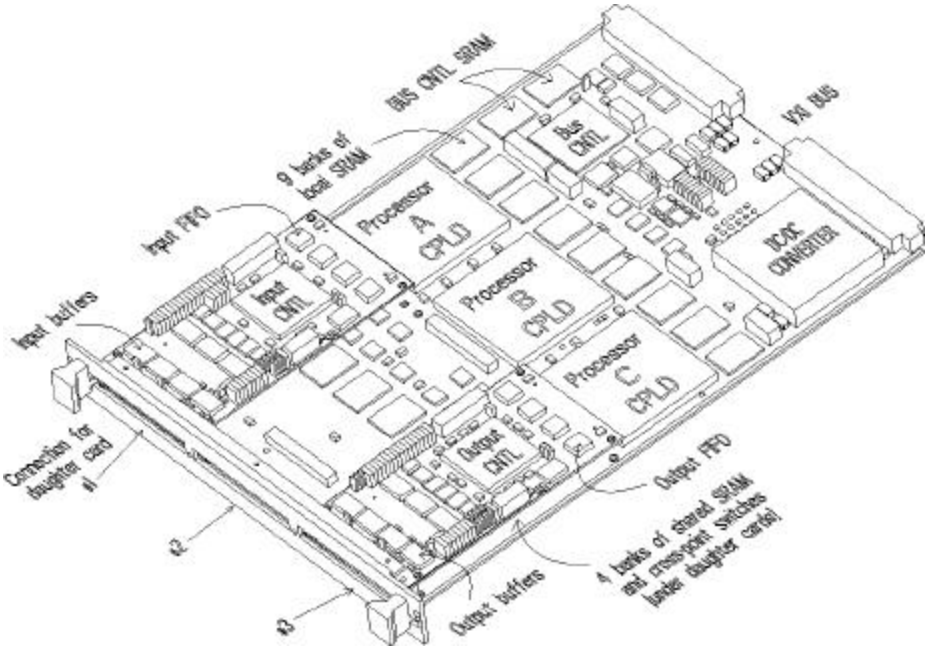

Fig. 1. Drawing of the RCA-2 board with an input and an output daughter card attached.

Figure #2 shows the functional blocks of the RCA-2 board and the data paths between them. Each functional block is described in more detail below.

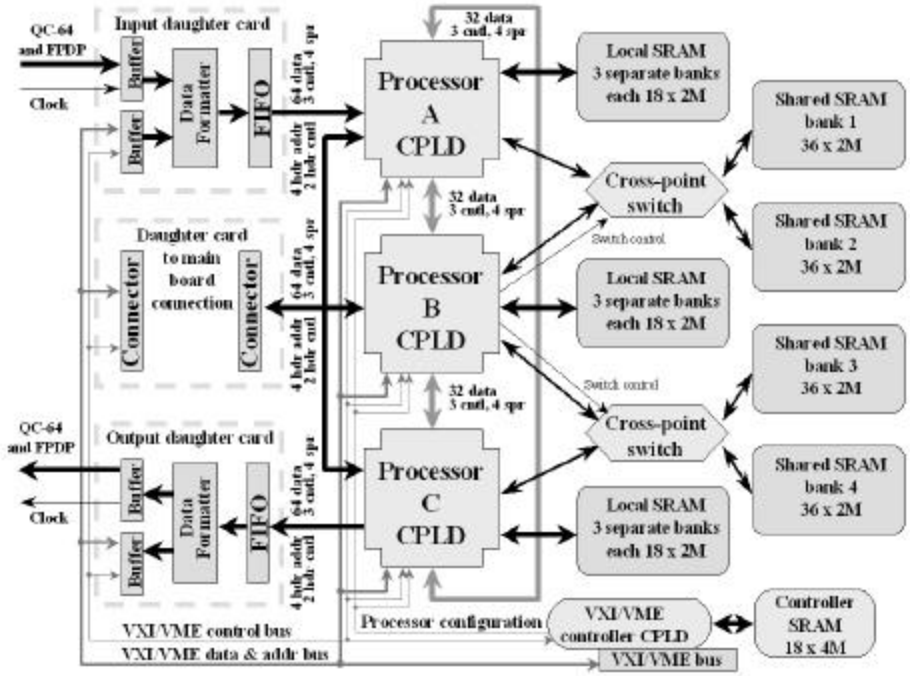

Fig. 2. Functional block diagram of the RCA-2 board and the input and output daughter cards.

3.1 Core processing CPLD

The processing core of the board is created with three Altera[2] 10K130V SRAM based Complex Programmable Logic Devices (CPLD's). The 10K130V implement Altera's architecture for the FLEX 10K series of parts (see the Altera data book for a full description of the architecture). The 10K130V have 6,656 logic elements each with a 4-input look-up table, a flip-flop and interconnects. Eight logic elements are grouped into logic array blocks for a total of 832. In addition to logic array blocks there are 16 embedded array blocks that provide 32Kbits of memory. Signals passing on and off the chip go through one of 464 user I/O cells. Signals between logic array blocks, embedded array blocks and I/O cells are routed with global interconnects that run the full length and width of the device.

[2] Altera Corporation, 101 Innovation Drive, San Jose, California USA 95134. http://www.altera.com.

The RCA-2 design team selected the Altera FLEX 10K series as the system processor for several reasons. First, to accomplish complex algorithms, the large amount of logic available in the 10K130V is required. Second, to accommodate the desired memory and data path connections, a large number of I/O pins are needed, which the Altera part supplied. Finally, with complex logic and lots of I/O connections, a large supply of long-line routing resources is needed to support designs. The global interconnects provided by Altera will be relied on to meet these routing requirements.

Power to the CPLD is provided by two separate sets of pins. One set of power pins, VCC_INT, supplies power to the internal core logic and requires 3.3 volts. (The 3.3 volt supply is created by an onboard DC/DC converter from the 24 volt supply available on the VXI backplane.) The second set of power pins, VCC_IO, supplies power to the I/O cells and can be driven by either 5 or 3.3 volts. Even when driven by 3.3 volts, the I/O cells are tolerant of 5 volt inputs. The RCA-2 board uses 3.3 volts for the I/O cells.

The RCA-2 board uses the 10K130V in the 599-pin Pin-Grid-Array (PGA) package. Altera will also be producing the 10K250A in a pin compatible package. The 10K250A nearly doubles the number of logic array blocks available and increases the number of embedded array blocks to 20. Altera has announced that the new 10K250E will be available in the same package. The FLEX 10KE series uses an advanced 0.25-micron, five-layer-metal CMOS SRAM process technology. This results in a smaller die, faster chip speeds and lower power consumption. However, the FLEX 10KE requires 2.5 volts for the internal core logic.

To accommodate both the 10K and 10KE series of parts, the RCA-2 board uses a separate power plane to supply VCC_INT. The VCC_INT plane supplies all the internal core logic power pins for all three CPLD processors and the system controller CPLD (discussed later). When using the current FLEX 10K series, the VCC_INT plane would be strapped to the 3.3 volt supply. When the new 10KE series parts are used the VCC_INT plane would need to be powered by a separate, 2.5 volt external supply. The separate VCC_INT plane will allow the performance of the board to be extended as the newer parts become available.

3.2 Local Processor Memory

Connected to each processor CPLD are three independent banks of synchronous static RAM. The RCA-2 board uses Zero Bus Turnaround (ZBT) SRAM from Micron[3] that can support either a read or a write operation on each clock cycle. Micron currently produces SRAM parts that have 18 bits of data and are 128K deep. They have announced plans to produce parts as large as 1M deep. Each bank of local memory has two chips for a maximum size of 18 bits by 2M. These parts can support clock rates of greater than 100 MHz.

Since each bank is independent, the processor can perform read or write operations on each bank simultaneously without bus contention. Additionally, because of its

[3] Micron Technology, Incorporated, 8000 S. Federal Way, Boise, Idaho USA 83707-0006. http://www.micron.com/mit.

synchronous operation, the memory fits cleanly with algorithm implementations inside the CPLD processor. This provides maximum flexibility to provide storage for intermediate results, for coefficients and can provide large circular buffers for delaying data.

3.3 Shared Memory

To provide a large data block transfer capability between pairs of processors, two banks of shared memory were created. These banks of shared memory are between processors A and B, and between processors B and C. All data, address and control signals between the processors and the shared memory go through a large cross-point switch. The cross-point switch was built using bus exchange switches from IDT[4]. This allows each processor to have complete control of the bank of shared memory that it is connected to without contention. Once one processor has written a complete data block, the two banks of memory are exchanged, making the data immediately available to the second processor. Then, while the second processor works on the new block of data, the first processor can start creating the next block of data.

Each bank of shared memory is created with 4 chips of ZBT SRAM for a maximum size of 36 bits by 2M. Exchanging the banks of memory is controlled by a signal from processor B. One entire bank of memory can be exchanged with the other bank in one clock period. The shared memory fully supports the pipeline processing of block oriented data by allowing one processor to exchange a complete data block with its neighboring processor.

3.4 Dedicated Data Paths

Processors A and C each have a dedicated connection to I/O daughter card location 1 and 3 respectively. These connections are 77-bits wide. The daughter card design and the configuration of the Altera processor determine how these bits are used. LANL has developed input and output cards (discussed later) that determine the functions of these buses.

Processor B also has a 77-bit connection to daughter card location 2, but it is not a dedicated connection. Instead, this bus is also connected to processors A and C. This shared bus has several uses. Of course it can provide a data path between the daughter card and processor B. However, the main use is to allow two input streams of data to be brought together in one processor. Or, alternatively, to allow two output streams of data to be produced by one processor. Since this bus is fully shared by processors A, B and C, the design provides the flexibility to steer the data where it needs to go. An additional use of this shared bus is for inter-processor communications if daughter card location 2 is not used for I/O.

[4] Integrated Device Technology Inc., 2975 Stender Way, Santa Clara, California USA 95054. http://www.idt.com.

Normal inter-processor communication takes place between processors over a 39-bit bus. The three processors are connected in a ring; A to B, B to C and C to A. As with the other data paths, the functions of these bits are determined by the configuration of the processors. LANL has adopted the initial convention that these buses would contain 32 bits of data, three control signals and four spare signals.

3.5 I/O Daughter Cards

Each of the three daughter card locations on the RCA-2 board provides two, 100-pin connectors. One connector is the data path connection from the daughter card to the Altera CPLD's plus ground connections. The second connector provides VXI bus access directly to the daughter card. Also included on the second connector are ground connections and various voltages from the backplane to power the daughter card. There is room at the edge of the daughter card for an 80-pin front panel data connector.

Although intended to provide a flexible front panel interface, daughter cards can fulfill other functions based on their design. One such function would be a high speed data path between a CPLD processor and the VXI bus. Another daughter card design might not implement any I/O functions at all but provide expanded memory resources instead. The design of a daughter card can be tailored to the requirements of the system and the algorithms being implemented.

For initial RCA-2 operations, LANL has developed an input daughter card and an output daughter card for high speed I/O. These daughter cards can implement both the Front Panel Data Port (FPDP) interface standard sanctioned by VITA[5] and the open QuickComm-64 (QC-64) standard from Catalina Research[6]. The cards provide First-In-First-Out (FIFO) memory for 64-bits of data going to or coming from the CPLD processor on the RCA-2 board. These FIFO's provide elastic buffers between cards, eliminating the requirement that all boards in the system must be synchronized and operate in lock step. LANL sees this as critically important for implementing block-oriented pipeline processing in a system containing multiple boards.

The QC-64 interface standard allows data to be transferred either as a continuous stream or as a packet. With packets of data, QC-64 allows header and trailer information to be appended to the data. All data transfers are synchronous with the front panel clock signal and are coordinated by three control signals. When transferring packets of data, two additional control signals are used to indicate when valid data begins and when it ends. These same signals control the transfer of header and trailer information. Header and trailer data is inserted in the place of normal sensor data and is further qualified by a 4-bit address indicating the function of the information. All data, header address and header control signals pass through the FIFO's.

---

[5] VMEbus International Trade Association, 7825 East Gelding Drive, Suite 104, Scottsdale, Arizona USA 85260. http://www.vita.com.

[6] Catalina Research Incorporated, 1321 Aeroplaza Drive, Colorado Springs, Colorado USA 80916. http://www.cri-dsp.com.

Based on the QC-64 interface standard, LANL has initially assigned functions to the 77-bit bus from the daughter card to the CPLD processors. These functions are: 64 bits of data, 3 data control signals, 4 bits of header address, 2 header control signals and 4 spare signals.

The function of the header and trailer information has not been fully defined. Since this information can travel with the data as it flows through the board, it will provide unique capabilities when implementing block-oriented pipeline processing. Information that might be included in the header could be gain, offset, start time and other sensor scaling information. For example, the current setting of an automatic gain control circuit could effect how a CPLD processor handles the data. Time stamps are especially valuable so that a processor can match together two blocks of data coming from separate sources. Trailer information might include an exponent value when implementing block floating-point representation.

3.6 VXI Interface

The RCA-2 board provides VXI bus access to each CPLD processor and to each daughter card. To each processor there are 8 bits of address and 18 bits of data from the backplane. The 18 bits for the data bus corresponds to the size of the local memory data bus. The VXI connection to the daughter cards includes 8 bits of address and a full 32 bits of data. The system controller CPLD supplies decoded VXI bus control signals and takes care of interrupts.

Through this VXI bus interface, data can be written to or read directly from any of the processors or daughter cards at standard VME bus rates. The main use of the VXI bus interface is for configuration, control and reporting functions. Examples of these functions include loading configuration files for the processors, loading coefficient values and monitoring the board's performance. These operations are expected to occur infrequently.

3.7 System Controller CPLD

The RCA-2 board uses an Altera 10K50V CPLD in a 356-pin Ball Grid Array (BGA) package to provide the VXI bus interface and to load configuration data into the processors. Altera has announced that a 10K100E will also be available in this package.

The system controller CPLD is automatically configured when power is turned on from onboard EPROM. Attached to the system controller is a local bank of ZBT SRAM. This bank has four chips for a maximum size of 18 bits by 4M. The controller also has a full VXI interface with 32 bits of data, 31 bits of address in addition to all data, address and interrupt control signals. These control signals are decoded and provided to the rest of the board. Depending on the configuration of the controller CPLD, the board can be either a memory mapped VME device or a VXI dynamically configured board.

The controller has separate passive serial configuration links with each of the board's CPLD processors. To load a processor, a configuration file is first written across the

backplane and stored in the controller's bank of SRAM. Then, when commanded, the controller reads the file and writes it out the requested configuration link. If the maximum amount of controller SRAM is installed, there is enough memory to store more than 32 different configuration files for the 10K130. This allows the files to be pre-loaded on the board so reconfiguration of a processor can be done as quickly as possible.

4 Theory of Operation

The RCA-2 board is primarily designed to support pipeline operations on blocks of data. In its most basic operation, the data streams onto the board from the front panel directly to a CPLD processor. This processor then operates on the data, using local memory as needed. After its last operation on the data, the processor then writes it to a bank of shared memory. When the entire block of data has been written, the banks of shared memory are exchanged and processing begins on the next block of data. Blocks of data move from processor to processor as required by the algorithm. After the third processor has finished, the data is written directly to the front panel.

Obviously there is tremendous flexibility to how data is handled and how algorithms are mapped to the RCA-2 board. For particular applications, passing data directly between processors may perform better than exchanging shared memory banks. Still other applications may alternate blocks of data between two processors and use the third processor to accumulate the results. These application related questions can only be answered by first understanding the requirements and the algorithms of a specific system.

A fundamental principle of the RCA-2 board design was to keep application data off the backplane bus or any other bus having more than one data transmitter. This principle led to the board's heavy use of dedicated data paths between processors and between the front panel data ports and processors. The board does have one data path with more than two components, that is the connection between daughter card #2 and all three CPLD processors. However, this data path allows two data streams to be joined or created by one processor as discussed above. For any given application, this shared data path should still comply with the 'only one data transmitter' rule.

By relying on dedicated data paths for moving data, the data throughput rate of a system is not dependent on the performance of the backplane bus. Therefore, additional RCA-2 boards can be added to a system without affecting data throughput. This would not be true if all the boards in a system relied strictly on the backplane to move data.

5 Processor Logic Development

Before the RCA-2 board can perform useful work, the CPLD processors must be properly configured. This requires a designer to translate signal processing algorithms into logic implemented in the Altera 10K130V devices. These algorithms may be initially presented to the designer as a series of mathematical formulas or as Matlab, DSP or 'C' code. It is then the designer's responsibility to translate the algorithms into schematics or to VHDL,

Verilog or some other RTL-level code. In addition to the algorithm, the designer must also incorporate data flow, memory interfaces, process coordination, system control, error handling and the other functions required to keep a system running correctly. All of this must then be partitioned across several CPLD processors and possibly across multiple boards. Finally, the logic must be placed and routed inside the CPLD's. In today's design environment, this is *not* a trivial effort.

To support logic development for the CPLD processors, LANL is developing a VHDL representation of the RCA-2 board, the two daughter cards and the VXI backplane. This representation will include behavior models for the ZBT memories, the FIFO memories, the cross-point switches and the VXI bus master CPU. To this testbench the designer would add the logic design of the three CPLD processors for simulation. This VHDL testbench will allow the designer to concentrate on developing his specific application code and to functionally test it at an early stage. A testbench of a complete system might include several RCA-2 boards, daughter cards and models of system specific boards all plugged into a backplane with front panel ports linked together.

LANL is also developing standard VHDL code at the RTL level that provides a basic VXI bus interface, a memory interface, daughter card interfaces, and allows configuration files to be loaded into the CPLD processors. Additionally, there is a package of code being developed that allows extensive testing of a board for diagnostic purposes. As application code is developed for the board, LANL intends to create a library of other useful routines.

6 RCA-2 Performance

At the time this paper is being written, no testing of the RCA-2 board has yet taken place since the board is currently being assembled. As test results become available, they will be posted to the project's web page at http://www.lanl.gov/rcc.

7 Conclusions

Los Alamos National Laboratory has specifically designed a reconfigurable computer board to support high speed processing of large blocks of data. Programmable logic is the key link between high speed, application specific, front-end sensors and slower speed general purpose computers. It is understood that the architecture of one reconfigurable processor, combined with its local memory, will not provide ideal implementation of all algorithms. However, by linking multiple processors and multiple boards together with high speed data paths, many algorithms can be implemented allowing an enormous increase in computing power and data throughput. Los Alamos National Laboratory believes that reconfigurable computing technology is a critical component of modern digital remote sensing systems.